arXiv:1603.00286v3 [cs.GT] 4 Aug 2016

How to Re-Divide a Cake Fairly *

Erel Segal-Halevi †

August 5, 2016

Abstract

This paper studies the classic problem of fairly dividing a heterogeneous resource ("cake") among several agents with different value measures. The innovation is that the resource is already divided among several "old" agents. When "young" agents come, the resource should be re-divided in a way that is fair both for the young and for the old agents.

Motivated by re-division of land and other real-life division problems, we present re-division protocols for various settings that differ in the geometric constraints on the cake and its pieces: (a) no geometric constraints - the cake is an arbitrary Borel subset of the plane and the pieces may be arbitrary Borel subsets of it; (b) connectivity - the cake is a one-dimensional interval and each piece must be a contiguous interval; (c) rectangularity - the cake is a two-dimensional rectangle or rectilinear polygon and the pieces should be rectangles; (d) convexity - the cake is a two-dimensional convex polygon and the pieces should be convex.

Our re-division protocols have implications on another problem: the price-of-fairness - the loss of social welfare caused by fairness requirements. Each protocol implies an upper bound on the price-of-fairness with the respective geometric constraints.

*This research was funded in part by the ISF grant 1083/13, the Doctoral Fellowships of Excellence Program and the Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University. I am grateful to Galya Segal-Halevi, Yonatan Aumann, Avinatan Hassidim, Reuven Cohen and Ioannis Caragiannis for helpful discussions and suggestions. I am also grateful to Math Stack Exchange members, particularly, Arseiny Akopyan, Chris Culter, Varun Dubey (vnd), Alex Ravsky, Swami Sarvattomananda (Shreesh), David K and JustPassingBy. All figures were created with GeoGebra 5 [33].

†Bar-Ilan University, ramat-gan 5290002, Israel. erelsgl@gmail.com

1 Introduction

Fair division of land and other resources among agents with different preferences has been an important issue since Biblical times. Today it is an active area of research in the interface of computer science [45, 46] and economics [42]. Its applications range from politics [14, 17] to multi-agent systems [20].

The classic setting assumes a one-shot division: the resource is divided once and for all, like a cake that is divided and eaten. But in practice, it is often required to re-divide an already-divided resource. One example is a cloud-computing environment, where new agents come and require resources held by other agents. A second example is fair allocation of radio spectrum among several broadcasting agencies: it may be required to re-divide the frequencies to accommodate new broadcasters. A third example is land-reform: large land-estates are held by a small number of landlords, and the government may want to re-divide them to landless citizens.

In the classic one-shot division setting, there are n agents with equal rights, and the goal is to give each agent a fair share of the cake. The common definition of a "fair share" is a piece worth at least $1/n$ of the total cake value, according to the agent's personal valuation function. This fairness requirement is usually termed *proportionality*. When proportionality cannot be attained, it is often (see Subsection 1.4) relaxed to *r-proportionality*, which means that each agent receives at least a fraction r/n of the total, where $r \in (0,1)$ is constant independent of n.

In contrast, in the re-division setting, there are $m < n$ "old" agents, who currently hold a piece of the cake, and $n - m$ "young" agents, and it may be required to give extra rights to the old agents. In particular, it may be required to give each old agent the opportunity to keep a substantial fraction of its current value. This may be due either to efficiency reasons (in the cloud computing scenario) or economic reasons (in the radio spectrum scenario) or political reasons (in the land-reform scenario). We call this requirement *retention*. Given a constant $r \in (0,1)$, *r-retention* means that each old agent receives at least r times its old value.

Can proportionality and retention be attained simultaneously? Our first result answers this question positively.

Theorem 1. *For every constant $r \in (0,1)$ and for every existing division of the cake, there exists a division that simultaneously satisfies r-proportionality and $(1-r)$-retention. Moreover, when r is a constant rational number, such a division can be found in time $O(n^2)$.*

The parameter r represents the level of compromise between the old and young agents: small r is better for the old and large r is better for the young. For example, taking $r = 1/2$, Theorem 1 implies that it is possible to re-divide the cake, giving each old agent at least half its previous value, while giving each agent (old or young) at least $1/(2n)$ of the total cake value.

1.1 Geometric constraints

While Theorem 1 is encouraging, it ignores an important aspect of practical division problems: geometry. The division it guarantees may be highly fractioned, giving each agent a large number of disconnected pieces. But in many practical division problems, the agents may want to receive a single connected piece. For example, when the divided resource is a time-interval, each agent may need a single contiguous interval rather than a large number of disconnected ones. Can partial-proportionality and partial-retention be attained simultaneously with a connectivity constraint? Our second result answers this question negatively.

Theorem 2. *When the cake is a 1-dimensional interval and each piece must be an interval, for every positive constants $r_1, r_2 \in (0,1)$, it may be impossible to simultaneously satisfy r_1-proportionality and r_2-retention. Moreover, if the cake should be re-divided such that each agent receives an interval with* any positive value*, then for every integer $k \in \{1,\ldots,m\}$, there might be k old agents who will receive at most a fraction $1/\lfloor \frac{n}{k} \rfloor$ of their old value.*

Note that the latter part of the theorem involves a property much weaker than proportionality: all we want is to guarantee each agent a positive value. With the connectivity constraint, even this weak "positivity" requirement is incompatible with r-retention for every constant $r > 0$: a positive division might require us to give one old agent at most $1/n$ of its previous value, give two old agents at most $2/n$ of their previous value, give $m/2$ old agents at most $\frac{1}{2}\frac{m}{n}$ of their previous value, etc.

Theorem 2 motivates the following weaker retention requirement: for every k, at least $m-k$ old agents receive at least a fraction $1/\lfloor\frac{n}{k}\rfloor$ of their old value. For example (taking $k = m/2$ and assuming all quotients are integers), at least half the old agents should receive at least $\frac{1}{2}\frac{m}{n}$ of their old value and (taking $k = m/4$) at least 75% of the old agents should receive at least $\frac{1}{4}\frac{m}{n}$ of their old value.

This criterion can be interpreted as follows. The quantity $\frac{m}{n}$ is the ratio of increase in population. We would like to let each old agent keep at least $\frac{m}{n}$ of its current value, so that the decrease in value is proportional to the increase in population. However, in light of Theorem 2, this guarantee cannot be made to all old agents (for example, there may be an old agent who receives only $1/n$ of his previous value). A constant-factor approximation of this guarantee can be made to some of the agents; the smaller the factor, the larger the number of agents that can get this guarantee. The definition is inspired by the "90th percentile" criterion common in Service-Level-Agreements and Quality-of-Service analysis, e.g. [23, 55]. It can also be justified by political reasoning: if a land-reform plan has to be approved by a majority of the old agents, then it may be sufficient to ensure that most old agents are happy, even if a small number of them are, inevitably, left unhappy.

Our third result almost matches the upper bound on retention implied by Theorem 2. The only difference is that the upper bound there is $1/\lfloor\frac{n}{k}\rfloor$ while the lower bound is $1/\lceil\frac{n}{k}\rceil$:

Theorem 3. *When the cake is a 1-dimensional interval and each piece must be an interval, it is possible to find in time $O(mn\log n)$ a 1/3-proportional division in which, for every integer $k \in \{1,\ldots,m\}$, at least $m-k$ old agents receive at least $1/\lceil\frac{n}{k}\rceil$ of their previous value.*

Theorem 3, like most cake-cutting papers, assumes that the cake is 1-dimensional. In realistic division scenarios, the cake is often 2-dimensional and the pieces should have a pre-specified geometric shape, such as a rectangle or a convex polygon. Rectangularity and convexity requirements are sensible when dividing land, exhibition space in museums, advertisement space in newspapers and even virtual space in web-pages. Moreover, in the frequency-range allocation problem, it is possible to allocate frequency ranges for a limited time-period; the frequency-time space is two-dimensional and it makes sense to require that the "pieces" are rectangles in this space [35].

2-dimensional cake-cutting introduces new challenges over the traditional 1-dimensional setting. As an example, In 1-D, it can be assumed that the initial allocation is a partition of the entire cake; this is without loss of generality, since any "blank" (unallocated part) can just be attached to a neighboring allocated interval without harming its shape or value. However, in 2-D, the initial allocation might contain blanks that cannot be attached to any allocated piece due to the rectangularity or convexity constraints. For example, suppose the cake is as the rectangle illustrated below:

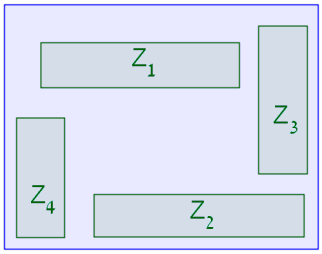

There are 4 old agents and each agent i has positive value-density only inside the rectangle Z_i. The most reasonable division (e.g. the only Pareto-efficient division) is to give each Z_i entirely to agent i. But, this allocation leaves a blank in the center of the cake, and this blank cannot be attached to any allocated piece due to the rectangularity constraint.

This counter-intuitive scenario cannot happen in a one-dimensional cake. Handling such cases requires new geometry-based tools. Using such tools we can prove analogues of Theorem 3 to two common 2-dimensional settings.

Theorem 4. *When the cake is a rectangle and each piece must be a parallel rectangle, it is possible to find in time $O(mn \log n)$ a 1/4-proportional division in which, for every integer $k \in \{1, \ldots, m\}$, at least $m - k$ old agents receive at least $1/\lceil \frac{n}{k} \rceil$ of their previous value.*

Theorem 5. *When the cake is a 2-dimensional convex polygon and each piece must be convex, it is possible to find in time $O(mn \log n)$ a 1/5-proportional division in which, for every integer $k \in \{1, \ldots, m\}$, at least $m - k$ old agents receive at least $1/\lceil \frac{n}{k} \rceil$ of their previous value.*

Remark. In the interval, rectangle and convex settings, the geometric constraints are mostly harmless without the retention requirement: when the cake is an interval/rectangle/convex, classic algorithms for proportional cake-cutting, such as Even-Paz [30], can be easily made to return interval/rectangle/convex pieces by ensuring that the cuts are parallel. Similarly, the retention requirement is easy to satisfy without the geometric constraints, as shown by Theorem 1. It is the combination of these two requirements that leads to interesting challenges.

Our next result generalizes Theorem 4 to a cake that is a *rectilinear polygon* - a polygon all whose angles are 90° or 270°. Rectilinearity is a common assumption in polygon partition problems [37]. The "complexity" of a rectilinear polygon is characterized by the number of its *reflex vertexes* - vertexes with a 270° angle. We denote the cake complexity by R. A rectangle - the simplest rectilinear polygon - has $R = 0$. The cake below has $R = 4$ reflex vertexes (circled):

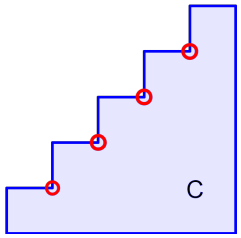

Theorem 6. *When the cake is a rectilinear polygon with R reflex vertexes, and each piece must be a rectangle, it is possible to find in time $O(mn \log n + poly(R))$ a division that guarantees each agent a fraction $1/(4n + R)$ of the total cake value, in which for every integer $k \in \{1, \ldots, m\}$, at least $m - k$ old agents receive at least $1/\lceil \frac{n}{k} \rceil$ of their previous value.* [1]

1.2 Application: bounding the price-of-fairness

Redivision protocols can be used not only to compromise between old and new agents, but also to compromise between fairness and efficiency. Often, the most economically-efficient allocation is not fair, while a fair allocation is not economically-efficient. The trade-off between fairness

[1] The guarantee of $1/(4n + R)$ is calculated as a fraction of the total cake value. However, with a rectilinear cake and a rectangular piece, even a single agent cannot always get the entire cake value to itself. Therefore, one could think of an alternative guarantee where the benchmark for each agent is the largest value that this agent can attain in a rectangle. For example, we could guarantee each agent a fraction $1/(4n)$ of the value of its most valuable rectangle. However, such guarantee might be much worse than the guarantee of Theorem 6. The proof in Appendix A implies that the value of the most valuable rectangle might be as small as $1/(R + 1)$ of the total cake value. Therefore, the alternative guarantee of $1/(4n)$ this value translates to a guarantee of $1/(O(n \cdot R))$ - much worse than the $1/(O(n + R))$ guaranteed by Theorem 6.

and efficiency is quantified by the *price-of-fairness* [3, 12, 13, 18]. It is defined as the worst-case ratio of the maximum attainable social-welfare to the maximum attainable social-welfare of a fair allocation. The social welfare is usually defined as the arithmetic mean of the agents' values (also called *utilitarian welfare*) or their geometric mean (also called *Nash welfare* [42]).

A redivision protocol can be used to calculate an upper bound on the price of fairness in the following way. Take a welfare-maximizing allocation as the initial allocation; use a redivision protocol to produce a partially-proportional allocation in which the utility of the old agents is close to their initial utility; conclude that the new welfare is close to the initial (maximal) welfare.

Without geometric constraints, we have the following upper bound:

Theorem 7. *For every $r \in [0, 1]$, the utilitarian-price of r-proportionality is at most $1/(1-r)$.*

Note that when $r = 1$, the bound is infinity. Indeed, [18] proved that the price of 1-proportionality in this setting is $\Theta(\sqrt{n})$, which is not bounded by any constant. Our results show that by making a small compromise on the level of proportionality we can get a constant (independent of n) bound on the utilitarian-price. The parameter r sets the level of trade-off between fairness and efficiency.

With geometric constraints, we have the following upper bounds:

Theorem 8. *When the cake is an interval and each piece must be an interval, for every $B \geq 3$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 8.4.*

Theorem 9. *When the cake is a rectangle and each piece must be a rectangle, for every $B \geq 4$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 11.2.*

Theorem 10. *When the cake is convex polygon and each piece must be convex, $\forall B \geq 5$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 14.*

Note that the first claim in Theorem 8 is subsumed by [3], who prove that the utilitarian-price of 1-proportionality in this setting is $\Theta(\sqrt{n})$. We bring this claim only for completeness. The second claim in this theorem, as well as the following theorems which deal with two-dimensional constraints, are not implied by previous results.

1.3 Paper layout

The introduction section is concluded by surveying some related work. Section 2 formally presents the model. Section 3 proves our theorem for unrestricted pieces. Section 4 proves our two one-dimensional theorems. Section 5 proves our three two-dimensional theorems. Section 6 proves our bounds on the price-of-fairness. Section 7 discusses some directions for future work.

Appendix A analyzes the problem of fairly dividing a rectilinear cake to rectangular pieces without retention requirements; it can be seen as a baseline for Theorem 6. Appendix B solves a problem in discrete geometry, which is used to analyze the redivision algorithm in the two-dimensional settings.

1.4 Related Work

Partial proportionality. While proportionality is the most common criterion of fair cake-cutting, it is often relaxed to partial-proportionality in order to achieve additional goals:

1. Speed: finding a proportional division takes $\Theta(n \log n)$ queries, but finding an r-proportional division takes only $\Theta(n)$ queries, for some positive but sufficiently small $r \leq 0.1$ [27, 28].

2. Improving the social welfare: proportional allocations may be socially inefficient; more socially-efficient allocations can be found by decreasing the value-guarantee per agent [2, 56].

3. Guaranteeing a minimum-size constraint: [19] consider a 1-dimensional setting in which each agent may get several intervals, but the length of each interval should be above some pre-specified threshold. They prove that it is impossible to guarantee an r-proportional allocation for any positive r, and provide some additive approximations.

4. Satisfying geometric constraints like square pieces [48, 49]. For example, when the cake is square and the pieces must be square, it is impossible to guarantee an r-proportional allocation for any $r \geq 1/2$, but there is an algorithm that guarantees a $1/4$-proportional allocation.

Dynamic fair division Our cake redivision problem differs from several division problems studied recently.

1. Dynamic resource allocation [31, 36] is a common problem in cloud-computing environments. The server has several resources, such as memory and disk-space. Agents (processes) come and depart. The server has to allocate the resources fairly among agents. When new agents come, the server may have to take some resources from existing agents. The goal is to do the re-allocation with minimal disruption to existing agents [31]. In this problem, the resources are *homogeneous*, which means that the only thing that matters is what quantity of each resource is given to each agent. In contrast, our cake is *heterogeneous* and different agents may have different valuations on it, so our protocol must decide which parts of the cake should be given to which agent.

2. Population monotonicity [41, 42, 51, 53] is an axiom that describes a desired property of allocation rules. When new agents arrive and the same division rule is re-activated, the value of all old agents should be weakly smaller than before. This axiom represents the virtue of solidarity: if sacrifices have to be made to support an additional agent, then everybody should contribute [51]. We, too, assume that old agents are taking part in supporting the new agents. However, we add the *retention* requirement, which means that old agents should be allowed to keep at least some of their previous value. In addition, while their approach is axiomatic and mainly interested in existence results, our approach is constructive and our goal is to provide an actual re-division protocol.

3. Private endowment, in economics resource allocation problems, means that each agent is endowed with a certain initial bundle of resources. Then, agents exchange resources using a certain market mechanism. The classic problem in economics involves homogeneous resources, but it has also been studied in the cake-cutting framework [5, 9]. A basic requirement in these works is *individual rationality*, which means that the final value allocated to each agent must be weakly *larger* than the value of the initial endowment (note the contrast with the population monotonicity axiom). In our problem we do not make this assumption as it is incompatible with fairness: since the young agents come with no initial endowment, individual rationality would mean that they might not receive anything in the exchange.

4. Online cake-cutting [54] is characteristic of a birthday party in an office, in which some agents come or leave early while others come or leave late. It is required to give some cake to agents who come early while keeping a fair share to those who come late. In contrast to our model, there it is impossible to re-divide allocated pieces, since they are eaten by their receivers. The fairness guarantees are inevitably weaker.

5. Land reform, the re-division of land among citizens, has been attempted in numerous countries around the globe and in many periods throughout history, from ancient Egypt in the

times of King Bakenranef (8th century BC) to modern-day Scotland (21st century AD). Fairness is a major concern in such reforms. Some books written on this topic are [11, 39, 44, 47].

Geometric cake models The most prominent cake-model is a one-dimensional interval, in which case the pieces are often required to be contiguous sub-intervals. This is especially true in computer science papers. Some exceptions (mostly from other disciplines) are:

1. The cake is a 1-dimensional circle and the pieces are contiguous arcs [6, 16, 52].

2. The cake is a 2-dimensional territory that lies among several countries. Each country should receive a piece adjacent to its border [7, 32].

3. The cake is 2-dimensional and the pieces are rectangles whose dimensions can be determined by the agents [35].

4. The cake is 2-dimensional and the pieces must be squares or fat polygons [48, 49].

5. The cake is 2-dimensional; the geometric constraints are connectivity, simple-connectivity or convexity [24].

6. The cake is multi-dimensional and the pieces should be simplexes or polytopes [10, 22, 34].

Many natural 2-dimensional settings have not been studied yet. For example, the setting studied here, where the cake is a rectilinear polygon and the pieces should be rectangles, has not been studied. As shown by the figure in the introduction, there is a qualitative (not only quantitative) difference between 2-D and 1-D division. 2-D division cake-cutting requires new techniques. For example, because we cannot assume that the entire cake is allocated, we must use a special *allocation-completion* step before our main redivision protocol (see Section 5) . The habit of assuming a one-dimensional cake might cause authors to miss such interesting steps. One goal of the present paper is to encourage authors to adapt a more realistic cake-model.

Price of fairness. The price-of-fairness has been studied in various contexts, such as routing and load-balancing [12, 13] and kidney exchange [25]. The price-of-fairness in cake-cutting has been studied in two settings:

- The cake is a *one-dimensional interval* and the pieces must be *intervals* [3]. The utilitarian-price-of-proportionality in this case is $\Theta(\sqrt{n})$.
- The cake is *arbitrary* and the pieces may be *arbitrary* [18]. The utilitarian-price-of-proportionality in this case is $\Theta(\sqrt{n})$ too.

Both papers study the price of other fairness criteria such as envy-freeness and equitability, but do not study the price in Nash-welfare. Additionally, they do not handle two-dimensional geometric constraints such as rectangularity or convexity.

Several authors study the algorithmic problem of finding a welfare-maximizing cake-allocation allocation in various settings:

1. The cake is an interval and the pieces must be connected [4];

2. The cake is an interval and the pieces must be connected, and additionally, the division must be proportional [8];

3. The cake and pieces are arbitrary, and the division must be envy-free [21].

4. The cake and pieces are arbitrary, and the division must be equitable [15].

2 Model

2.1 Cake division

The *cake* C is a Borel subset of a d-dimensional Euclidean plane $\mathbb{R}^d$. In this paper we focus on the common cases in which $d = 1$ and C is an interval, or $d = 2$ and C is a polygon. A *piece* is a Borel subset of C.

C has to be divided among $n \geq 1$ *agents*. Each agent $i \in \{1, \ldots, n\}$ has a *value-density* function v_i, which is an integrable, non-negative and bounded function on C. The *value* of a piece X_i to agent i is marked by $V_i(X_i)$ and it is the integral of its value-density:

$$V_i(X_i) = \int_{x \in X_i} v_i(x)dx$$

The V_i are measures and are *absolutely continuous with respect to the Lebesgue measure*, i.e., any piece with zero area has zero value to all agents. Therefore, we do not need to worry about who gets the boundary of a piece, since its value is 0.

The division protocols access the value measures via *queries* [46]: an *eval query* asks an agent to reveal its value for a specified piece of cake; a *mark query* asks an agent to mark a piece of cake with a specified value. The present paper ignores strategic considerations and assumes that agents answer truthfully.

The geometric constraints, if any, are represented by a pre-specified family S of *usable pieces*. In this paper, S will either be the set of all pieces (which means that there are no geometric constraints), or the set of all intervals, or the set of all rectangles, or the set of all convex pieces. We assume that each agent can use only a single piece from the family S.

An *allocation* is a vector of n pieces, $X = (X_1, \ldots, X_n)$, one piece per agent, such that the X_i are pairwise-disjoint and $\cup_{i=1}^{n} X_i \subseteq C$. Note that some cake may remain unallocated, i.e, *free disposal* is assumed. We explained in the introduction why this may be important. An *S-allocation* is an allocation in which all pieces are usable, i.e, $\forall i : X_i \in S$.

For every constant $r \in (0, 1)$, an allocation X is called *r-proportional* if every agent receives at least r/n of the total cake value:

$$\forall i \in \{1, \ldots, n\} : V_i(X_i) \geq \frac{r}{n} \cdot V_i(C)$$

A 1-proportional division is usually called in the literature "proportional". We also define a *positive allocation* as an allocation in which each agent receives a positive value: $\forall i : V_i(X_i) > 0$.

2.2 Cake redivision

The n agents are classified into two classes: m agents (indexed $1, \ldots, m$) are called *old* and the other $n - m$ agents (indexed $m + 1, \ldots, n$) are called *young*. There is an existing S-allocation of the cake C to the m old agents: $Z_1, \ldots, Z_m$. It is assumed that the old pieces Z_j are pairwise-disjoint and $\forall j : Z_j \in S$, but nothing else is assumed on the division. In particular, the initial division is not necessarily proportional, and some of C may be undivided.

It is required to create a new S-allocation of C to all agents: $X_1, \ldots, X_n$. For every constant $r \in (0, 1)$, the re-allocation satisfies the *r-retention* property if every old agent receives at least a fraction r of its old value:

$$\forall j \in \{1, \ldots, m\} : \qquad V_j(X_j) \geq r \cdot V_j(Z_j)$$

Since r-retention is not always compatible with partial-proportionality, we define the following weaker property. A re-allocation satisfies the *gradual-retention* property if, for every $k \in$

$\{1,\ldots,m\}$, there are at least $m-k$ indexes $j \in \{1,\ldots,m\}$ for which:

$$V_j(X_j) \geq \frac{1}{\lceil n/k \rceil} \cdot V_j(Z_j).$$

2.3 Social-welfare and Price-of-fairness

In addition to fairness, it is often required that a division has a high *social welfare*. The social welfare of an allocation is a certain aggregate function of the normalized values of the agents (the normalized value is the piece value divided by the total cake value). Common social welfare functions are sum (utilitarian) and product (Nash) [42]. We normalize them such that the maximum welfare is 1:

- *Utilitarian welfare* - the arithmetic mean of the agents' normalized values:

$$W_{util}(X) = \frac{1}{n} \sum_{i \in \{1,\ldots,n\}} \frac{V_i(X_i)}{V_i(C)}$$

- *Nash welfare* - the geometric mean of the agents' normalized values:

$$W_{Nash}(X) = \left(\prod_{i \in \{1,\ldots,n\}} \frac{V_i(X_i)}{V_i(C)} \right)^{1/n}$$

The goal of maximizing the social welfare is not always compatible with the goal of guaranteeing a fair share to every agent. For example, [18] describe a simple example in which the maximum utilitarian welfare of a proportional allocation is $O(1/n)$ while the maximum utilitarian welfare of an arbitrary (unfair) allocation is $O(1/\sqrt{n})$. This means that society has to pay a price, in terms of social-welfare, for insisting on fairness. This is called the *price of fairness*. Formally, given a social welfare function W and a fairness criterion F, the price-of-fairness relative to W and F (also called: "the W-price-of-F") is the ratio:

$$\frac{\max_X W(X)}{\max_{Y \in F} W(Y)} \qquad (*)$$

where the maximum at the nominator is over all allocations X and the maximum at the denominator is over all allocations Y that also satisfy the fairness criterion F. The cited example shows that the utilitarian-price-of-proportionality might be $\Omega(\sqrt{n})$.

When there are geometric constraints, they affect both the numerator and the denominator of (*). Therefore, it is not a-priori clear whether the price-of-fairness with constraints is higher or lower than without constraints.

3 Arbitrary Cake and Arbitrary Pieces

This section proves Theorem 1, which assumes no geometric constraints on the cake or pieces. Our main lemma is:

Lemma 3.1. *Given cake-allocations Z and Y and a constant $r \in [0,1]$, there exists an allocation X such that, for every agent i:*

$$V_i(X_i) \geq rV_i(Y_i) + (1-r)V_i(Z_i)$$

Moreover, when r is a constant rational number, X can be found using $O(n^2)$ mark/eval queries.

Proof. The space of utilities of cake-allocations is convex [26]. Therefore, there exists an allocation X satisfying the requirement as an equality: $\forall i : V_i(X_i) = rV_i(Y_i) + (1-r)V_i(Z_i)$.

Since the Dubins-Spanier theorem [26] is not constructive, we give here a constructive protocol for creating the allocation Z when r is a rational number, $r = p/q$ with $p < q$ some positive integers. For every pair of agents i, j (including when $i = j$), the protocol works as follows:

> **Step 1.** Agent i divides the piece $Z_i \cap Y_j$ to q pieces that are equal in its eyes.
> **Step 2.** Agent j takes the p pieces that are best in its eyes.
> **Step 3.** Agent i takes the remaining $q-p$ pieces.
> (Note that when $i = j$, agent i receives the entire piece $Z_i \cap Y_i$).

Each agent i is allocated a piece X_i which is a union of nq pieces: np pieces that agent i took from other agents (including itself) in piece Y_i and $n(q-p)$ pieces that were left for agent i from other agents in piece Z_i.

From every piece $Y_i \cap Z_j$ (for $j \in \{1, \ldots, n\}$), agent i picks the best p out of q pieces, which give it a value of at least $\frac{p}{q}V_i(Y_i \cap Z_j)$. Its total value of these np pieces is thus at least $rV_i(Y_i)$.

In addition, from every piece $Z_i \cap Y_j$ (for $j \in \{1, \ldots, n\}$), agent i receives $q-p$ out of q equal pieces, which give it a value of exactly $\frac{q-p}{q}V_i(Z_i \cap Y_j)$. Its total value of these $n(q-p)$ pieces is thus exactly $(1-r)V_i(Z_i)$. □

Proof of Theorem 1. Apply Lemma 3.1, with:

Y - any proportional allocation, which can be found by classic protocols such as [30, 50].

Z - the initial allocation to the old agents, where the young agents have empty pieces. □

Note that the redivision protocol gives to each old agent a piece that is not only worth at least $(1-r)V_i(Z_i)$, but is also a subset of Z_i (in addition to a subset of Y_i). This may be desirable in some cases, e.g. in land division, the old landlords may care not only about their value but also about getting a subset of their old plot.

Remark. The $O(n^2)$ complexity assumes the integers p, q are constant (not part of the input). If they are considered part of the input, then the complexity becomes linear in q which is exponential in the number of input bits. The number of queries can be reduced using concepts from number theory, but this is beyond the scope of this paper. See [40, 46] for details.

4 Interval Cake and Interval Pieces

In this section, the cake is an interval in $\mathbb{R}$. Each piece in the initial division is an interval in C and each piece in the new division must be an interval in C. We begin by proving the impossibility result (Theorem 2), using a lemma opposite to Lemma 3.1.

Lemma 4.1. *Let Z be a connected allocation and $k \leq n$ an integer. Then there exist valuations such that, in every* positive *connected allocation X, for every agent $j \in \{1,\ldots,k\}$: $V_j(X_j) \leq V_j(Z_i)/\lfloor \frac{n}{k} \rfloor$.*

Proof. Assume that the valuations are as follows. Each agent $j \in \{1,\ldots,k\}$ values the piece Z_j as $\lfloor \frac{n}{k} \rfloor$ and the rest of the cake as 0. The value-density of j in Z_j is piecewise-uniform: It has $\lfloor \frac{n}{k} \rfloor$ regions with a value of 1 and $\lfloor \frac{n}{k} \rfloor - 1$ "gaps" - regions with a value of 0, like this:

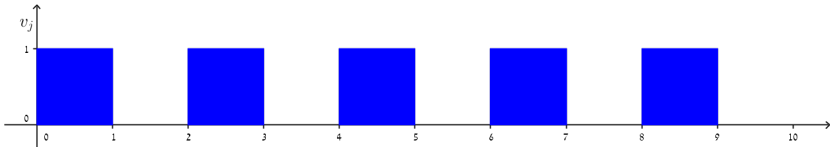

The other $n-k$ agents are divided to k groups of roughly equal size: the size of each group is either $\lfloor \frac{n-k}{k} \rfloor = \lfloor \frac{n}{k} \rfloor - 1$ or $\lceil \frac{n-k}{k} \rceil = \lceil \frac{n}{k} \rceil - 1$. Each agent in group j assigns a positive value only to a unique gap in the piece Z_j (so when the group size is $\lfloor \frac{n}{k} \rfloor - 1$, each gap is wanted by exactly one agent; otherwise, there is one gap wanted by two agents). In any positive division, each gap in Z_j must be at least partially allocated to an agent in group j. Hence, the interval allocated to agent j must contain at most a single positive region in Z_j - it is not allowed to overlap any gap. Therefore the value of agent j is at most $V_j(Z_j)/\lfloor \frac{n}{k} \rfloor$. □

Proof of Theorem 2. Apply Lemma 4.1 with Z = the initial allocation. □

To prove the matching positive result (Theorem 3), we introduce a protocol for fair division of an "archipelago" - a cake made of one or more interval "islands".

Lemma 4.2. *Let C be a cake made of $m \geq 1$ pairwise-disjoint intervals: $C = Z_1 \cup \cdots \cup Z_m$. There exists a division X of C among n agents, in which (a) Each agent i receives an interval entirely contained in one of the islands: $\forall i : \exists j : X_i \subseteq Z_j$, and (b) Each agent receives a value of at least $V_i(C)/(n+m-1)$. Moreover, X can be found using $O(mn \log n)$ mark/eval queries.*

Proof. We normalize the value measures of all agents such that the total value of C is $n+m-1$. The following recursive protocol allocates each agent an interval with a value of at least 1.

Base: $m = 1$. The cake is a single interval and its total value is n. Use the Even-Paz protocol [30] to allocate each agent an interval with a value of at least 1.
Step: $m > 1$.

1. Ask each agent $i \in \{1,\ldots,n\}$ to evaluate the island Z_m.
2. Order the agents in descending order of their evaluation: $V_1(Z_m) \geq \cdots \geq V_n(Z_m)$.
3. Let q be the largest integer such that $V_q(Z_m) \geq q$ (or 0 is already $V_1(Z_m) < 1$).
4. If $q = 0$, discard the island Z_m. Otherwise ($q \geq 1$),
   divide Z_m proportionally among the agents $\{1,\ldots,q\}$ using Even-Paz protocol [30].
5. Divide the remaining $m-1$ islands recursively among the remaining $n-q$ agents.

The descending order of the agents guarantees that: $V_1(Z_m) \geq \cdots \geq V_q(Z_m) \geq q$. So in step #4, the interval Z_m is divided proportionally among q agents that value it as at least q, and each of these agents receives an interval with a value of at least 1.

By definition of q, $V_{q+1}(Z_m) < q+1$ (this is true even when $q = 0$). By the descending order of the agents, the same is true for all remaining agents $\{q+1,\ldots,n\}$. Therefore, all remaining agents value the remaining cake as more than $(m+n-1)-(q+1) = (n-q)+(m-1)-1$. Since there are $n-q$ agents and $m-1$ islands, the recursive algorithm gives each agent an interval with value at least 1. The Even-Paz protocol requires $O(n \log n)$ queries, and it is done at most m times, so the total number of queries is $O(mn \log n)$. □

Remark. The fraction of $1/(n+m-1)$, guaranteed by Lemma 4.2, is the largest that can be guaranteed. To see this, assume that all agents $i \in \{1,\ldots,n\}$ have the same value-measures - they value the islands $Z_1,\ldots,Z_{m-1}$ as 1 and the island Z_m as n (so their total cake value is $n+m-1$). The piece of every agent must be entirely contained in a single island. If any agent receives a piece in islands $Z_1,\ldots,Z_{m-1}$, then that agent receives a value of at most 1. Otherwise, if all n agents receive a piece in Z_m, then the value of at least one agent is at most 1. In both cases, at least one agent receives a fraction of at most $1/(n+m-1)$ the cake value. □

Proof of Theorem 3. Our protocol for re-division of an interval has three steps.

Step 1. Given the original partial allocation $Z_1 \cup \cdots Z_m \subseteq C$, extend it to a complete allocation $Z'_1 \cup \cdots Z'_m = C$, by attaching each "blank" (unallocated interval in C) arbitrarily to one of the two adjacent allocated intervals. This, of course, does not harm the values of the old agents: $\forall j \in \{1,\ldots,m\} : V_j(Z'_j) \geq V_j(Z_j)$.

Step 2. For each old agent $j \in \{1,\ldots,m\}$, add a "virtual agent" j^* and assign it a value-density function v^*_j:

$$v^*_j(x) = v_j(x) \qquad \text{if } x \in Z'_j$$
$$v^*_j(x) = 0 \qquad \text{if } x \notin Z'_j$$

Use the protocol of Lemma 4.2 on the $n+m$ agents, regarding the cake C as an archipelago and the pieces $Z'_1,\ldots,Z'_m$ as the islands.

Step 3. Give each old agent $j \in \{1,\ldots,m\}$ either the interval allocated to its normal agent j or the interval allocated to its virtual agent j^*, whichever is more valuable for it.

We now prove that the resulting allocation is 1/3-proportional and satisfies the gradual-retention property.

(a) *Proof of 1/3-proportionality.* By Lemma 4.2, each of the $n+m$ agents receives an interval contained in one of the pieces $Z'_1,\ldots,Z'_m$, with a value of at least $1/((n+m)+m-1)$ its total cake value. Since $m \leq n$, this value is larger than $1/(3n)$.

(b) *Proof of gradual retention.* We focus on the m virtual agents. First, by Lemma 4.2, every virtual agent j^* must receive an interval contained in Z'_j, since its value is positive only in the island Z_j. Moreover, by the pigeonhole principle, for every integer $k \leq m$, at most k islands are populated by at least $\lceil \frac{n}{k} \rceil$ normal agents. Hence, at least $m-k$ islands are populated by at most $\lceil \frac{n}{k} \rceil - 1$ normal agents. Together with the virtual agent, these islands are populated by at most $\lceil \frac{n}{k} \rceil$ agents. Hence, the proportional allocation of step #4 in the protocol of Lemma 4.2 gives these virtual agents an interval subset of Z'_j, which is worth for agent j at least $V_j(Z'_j)/\lceil \frac{n}{k} \rceil$. □

5 Polygonal Cake and Polygonal Pieces

Rectangle cake and pieces. Initially, we assume that the cake is a rectangle in $\mathbb{R}^2$. Each piece in the initial division is a rectangle parallel to C and each piece in the new division must be a rectangle parallel to C.

We would like to use the re-division protocol of Theorem 3. Steps #2 and #3 are easily adapted: the Even-Paz protocol [30] can operate on a rectangular cake, requiring the agents to make cuts parallel to the cake sides. This guarantees that the pieces are rectangles.

However, there is one obstacle. Step #1, the allocation-completion step, is no longer trivial. We cannot just attach each unallocated part of C to an allocated rectangle, since the result will not necessarily be a rectangle. We still need to extend the initial partial allocation $Z_1 \cup \cdots Z_m \subseteq C$ to a complete allocation, but the number of rectangles in the complete allocation might be larger than m, since we might have unattached blanks.

Our goal, then, is to find a partition of C to rectangles, $Z'_1 \cup \cdots Z'_{m+b} = C$, with $b \geq 0$, such that every input rectangle is contained in a unique output rectangle: $\forall j \in \{1, \ldots, m\} : Z_j \subseteq Z'_j$. The additional b rectangles are called blanks. In Step 3, the value guarantee per agent will be $1/((n+m)+(m+b)-1) = 1/(n+2m+b-1)$; therefore, we would like the number of blanks b to be as small as possible.

An example of the input and output of the allocation-completion step is shown below. Here, $b = 1$ since there is one blank - Z'_5. In this case $b = 1$ is minimal.

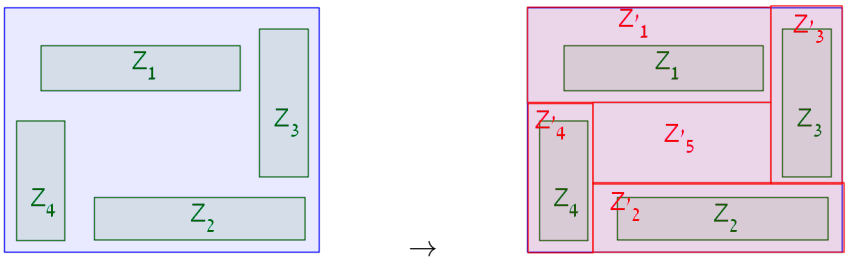

Convex cake and pieces. The situation is similar when C is convex and the pieces should be convex. The Even-Paz protocol can operate on a convex cake, requiring the agents to make cuts parallel to the each other. This guarantees that the pieces will be convex. In Step #1, a similar challenge arises. We have an initial partial allocation $Z_1 \cup \cdots Z_m \subseteq C$, where each Z_j is convex. We need a complete allocation $Z'_1 \cup \cdots Z'_{m+b} = C$, where each Z'_j is convex, every input piece is contained in a unique output piece, and the number of blanks b is minimal.

Rectilinear cake and rectangular pieces. There are efficient algorithms for partitioning a rectilinear polygon to a minimal number of rectangles. A rectilinear polygon with R reflex vertexes can be partitioned in time $O(poly(R))$ to at most $R+1$ rectangles [29, 37], and this bound is tight when the vertexes of C are in general position. Our goal is to bound b - the number of blank *rectangles*. Therefore, it is expected that the bound should depend on R, in addition to m.

The allocation-completion step for all two-dimensional settings is handled in Appendix B. While the allocation-completion algorithm is simple, its analysis is quite challenging and requires tools from discrete geometry. In Appendix B we prove:

Lemma 5.1. *There is an $O(m)$-time algorithm that extends a partial allocation $Z_1 \cup \cdots Z_m \subseteq C$ to a complete allocation $Z'_1 \cup \cdots Z'_{m+b} = C$, such that there are:*

(a) at most $m-3$ rectangular blanks when the cake and pieces are parallel rectangles:

(b) at most $2m-5$ convex blanks when the cake and pieces are convex polygons;

(c) at most $m+R-3$ rectangular blanks when the cake is rectilinear with R reflex vertexes and the pieces are rectangles. In this case the run-time is $O(m+poly(R))$.

The numbers of blanks in all cases are asymptotically tight.

Proof of Theorems 4, 5 ,6 . Use the protocol in the proof of Theorem 3, plugging into Step #1 the algorithm of Lemma 5.1. The value per agent is at least $1/(n+2m+b-1)$, which is:

(a) at least $1/(n+3m-4) > 1/(4n)$ in the rectangle case - satisfying 1/4-proportionality;

(b) at least $1/(n+4m-6) \geq 1/(5n)$ in the convex case - satisfying 1/5-proportionality;

(c) at least $1/(n+3m+R-4) > 1/(4n+R)$ in the rectilinear case. □

6 Price-of-Fairness Bounds

In this section, our redivision protocols are used to prove upper bounds on the price of partial-proportionality.

Theorem 7 follows directly from Theorem 1 by taking the original division to be a utilitarian-optimal division.

The proofs of Theorems 8, 9 and 10 are similar; only the constants are different. We present below only the proof of Theorem 9; to get the proofs of the other theorems, replace the constant "4" with "3" or "5" respectively.

The first part of Theorem 9 - regarding the **utilitarian** price - is proved by the following:

Lemma 6.1. *Let Z be a utilitarian-optimal rectangular division of a cake C among n agents who value the entire cake C as 1. Let U be the utilitarian welfare of Z:*

$$U := \frac{1}{n}\sum_{j=1}^{n} V_j(Z_j)$$

Then, there exists a $(1/4)$-proportional rectangular allocation of C to these same n agents with utilitarian welfare W, such that $U/W \in O(n^{1/2})$.

Proof. Apply the redivision protocol of Section 5 to the existing division by setting $m = n$ and treating all n agents as "old". The partial-proportionality guarantee of that protocol ensures that the new division is 1/4-proportional. The partial-retention of that protocol ensures that for every integer $k \in \{0, \ldots, n\}$, there is a set S_k containing at least $n - k$ agents whose value is more than $\max(\frac{kV_j(Z_j)}{2n}, \frac{1}{4n})$. Renumber the agents in the following way. Pick an agent from S_{n-1} (which contains at least one agent) and number it $n-1$. Pick an agent from S_{n-2} (which contains at least one other agent) and number it $n-2$. Continue this way to number the agents by $k = n-1, \ldots, 0$. Now, the utilitarian welfare of the new division is lower-bounded by:

$$W > \frac{1}{n}\sum_{k=0}^{n-1} \max(\frac{kV_k(Z_k)}{2n}, \frac{1}{4n}) \geq \frac{1}{n} \cdot \frac{1}{4n} \cdot \sum_{k=0}^{n-1} \max(kV_k(Z_k), 1)$$

and the utilitarian welfare ratio is at most:

$$\frac{U}{W} < 4n \cdot \frac{\sum_{k=0}^{n-1} V_k(Z_k)}{\sum_{k=0}^{n-1} \max(k \cdot V_k(Z_k), 1)}$$

Denote the ratio in the right-hand side by $\frac{\mathbf{NUM}}{\mathbf{DEN}}$. Let $a_k = V_k(Z_k)$, so that $\mathbf{NUM} = \sum_{k=0}^{n-1} a_k$ and $\mathbf{DEN} = \sum_{k=0}^{n-1} \max(k \cdot a_k, 1)$. To get an upper bound on U/W, we find a sequence $a_0, \ldots, a_{n-1}$ that maximizes $\frac{\mathbf{NUM}}{\mathbf{DEN}}$ subject to $\forall k : 0 \leq a_k \leq 1$.

Observation 1. in a maximizing sequence, $a_0 = 1$ and there is no $k > 0$ such that $a_k < 1/k$. *Proof*: Setting such a_k to $1/k$ increases **NUM** and does not change **DEN**.

Observation 2. A maximizing sequence must be weakly-decreasing (for all $k < k'$, $a_{k'} \geq a_k$). *Proof*: if there exists $k < k'$ such that $a_k < a_{k'}$, then we can swap a_k with $a_{k'}$. This does not change **NUM** but strictly decreases **DEN**.

Observation 3. In a maximizing sequence, there is no $k > 0$ such that $1/k < a_k < 1$. *Proof:*[2] If $1/k < a_k < 1$ then for some sufficiently small $\epsilon > 0$, both $a_k + \epsilon$ and $a_k - \epsilon$ are in $(1/k, 1)$ and replacing a_k with $a_k \pm \epsilon$ makes the ratio strictly smaller than the maximum. Replacing a_k with $a_k + \epsilon$ makes the ratio $\frac{\mathbf{NUM}+\epsilon}{\mathbf{DEN}+k\epsilon}$; this new ratio is smaller than $\frac{\mathbf{NUM}}{\mathbf{DEN}}$ so $\epsilon \cdot \mathbf{DEN} < k\epsilon \cdot \mathbf{NUM}$. Replacing a_k with $a_k - \epsilon$ makes the ratio $\frac{\mathbf{NUM}-\epsilon}{\mathbf{DEN}-k\epsilon}$; that new ratio is smaller

[2] We are grateful to Varun Dubey for suggesting this proof in: http://math.stackexchange.com/q/1609071/29780

than $\frac{\mathbf{NUM}}{\mathbf{DEN}}$ so $-\epsilon \cdot \mathbf{DEN} < -k\epsilon \cdot \mathbf{NUM}$. But the two latter inequalities $\epsilon \cdot \mathbf{DEN} < k\epsilon \cdot \mathbf{NUM}$ and $-\epsilon \cdot \mathbf{DEN} < -k\epsilon \cdot \mathbf{NUM}$ are contradictory. Hence, the assumption $1/k < a_k < 1$ is false.

Observations 1-3 imply that a maximizing sequence has a very specific format. It is characterized by an integer $l \in \{0, \ldots, n-1\}$ such that, for all $k \leq l$, $a_k = 1$ and for all $k \geq l+1$, $a_k = 1/k$. So:

$$\frac{\mathbf{NUM}}{\mathbf{DEN}} = \frac{\sum_{k=0}^{n-1} a_k}{\sum_{k=0}^{n-1} \max(k \cdot a_k, 1)} = \frac{(l+1) + (H_{n-1} - H_l)}{\frac{1}{2}l(l+1) + (n-l-1)} < \frac{2(l + H_n + 1)}{l^2 - l + 2(n-1)}$$

where $H_n = \sum_{k=1}^{n}(1/k)$ is the n-th harmonic number.

The number l is integer, but the expression is bounded by the maximum attained when l is allowed to be real. By standard calculus we get that the real value of l which maximizes the above expression is $l = \sqrt{2(n-1) + (H_n+1)(H_n+2)} - (H_n+1) = \Theta(\sqrt{n})$. Substituting into the above inequality gives:

$$\frac{\mathbf{NUM}}{\mathbf{DEN}} \leq \frac{\Theta(n^{1/2})}{\Theta(n)} = \Theta(n^{-1/2}) \qquad \Longrightarrow \qquad \frac{U}{W} < 4n \cdot \frac{\mathbf{NUM}}{\mathbf{DEN}} = O(n^{1/2})$$

as claimed. □

The second part of Theorem 9 - regarding the **Nash** price - is proved by the following:

Lemma 6.2. *Let Z be a Nash-optimal rectangular division of a cake C among n agents who value the entire cake C as 1. Let U be the Nash welfare of Z (the geometric mean of the values):*

$$U^n = \prod_{j=1}^{n} V_j(Z_j)$$

Then, there exists a $(1/4)$-proportional rectangular allocation of C to these same n agents with Nash welfare W, and $U/W < 11.2$.

Proof. Apply the redivision protocol of Section 5 to redivide the existing n pieces among the n agents. Renumber the agents as in Lemma 6.1. The Nash welfare of the new division, raised to the n-th power, can be bounded as:

$$W^n > \prod_{k=0}^{n-1} \max(\frac{k \cdot V_k(Z_k)}{2n}, \frac{1}{4n}) \geq (\frac{1}{4n})^n \prod_{k=0}^{n-1} \max(k \cdot V_k(Z_k), 1)$$

and the ratio of the new welfare to the previous welfare can be bounded as:

$$\frac{U^n}{W^n} < (4n)^n \cdot \frac{\prod_{k=0}^{n-1} V_k(Z_k)}{\prod_{k=0}^{n-1} \max(kV_k(Z_k), 1)} = \frac{(4n)^n}{\prod_{k=0}^{n-1} \max(k, 1/V_k(Z_k))}$$

The nominator does not depend on the valuations, so the ratio is maximized when the denominator is minimized. This happens when each factor in the product is minimized. The minimal value of the 0-th factor is 1 and the minimal value of the other factors is k. Hence:

$$\frac{U^n}{W^n} < \frac{(4n)^n}{\prod_{k=1}^{n-1} k} = \frac{(4n)^n}{(n-1)!} = \frac{n(4n)^n}{n!} \approx \frac{n(4n)^n}{\sqrt{2\pi n}(n/e)^n} = \sqrt{\frac{n}{2\pi}} \cdot (4e)^n$$

where e is the base of the natural logarithm. Taking the n-th root gives $U/W < (4e) \cdot \sqrt{n/2\pi}^{1/n}$. A calculation in Wolfram Alpha shows that the rightmost term $\sqrt{n/2\pi}^{1/n}$ is bounded globally by 1.03, so all in all $U/W < 1.03 \cdot 4 \cdot e < 11.2$ as claimed. □

7 Future Work

Two-dimensional division, the re-division problem and the price-of-fairness are relatively new topics, and there is a lot of room for future research in each of them.

7.1 Handling other geometric constraints

Two steps in our redivision algorithm are sensitive to the geometric constraint: the *allocation-completion algorithm* (Step #1 in Theorem 3), and the *Even-Paz protocol* (Step #4 in Lemma 4.2). We describe how these steps are affected by several alternative constraints.

1. Convexity in three or more dimensions. The Even-Paz protocol can easily operate on a multi-dimensional convex object, requiring the agents to cut using hyper-planes parallel to each other. However, we currently do not have an allocation-completion algorithm for convex objects (or even for boxes) in three or more dimensions.

2. Path-connectivity in two dimensions. If the pieces have to be path-connected but not necessarily convex, then the allocation-completion step is much easier and no blanks are created - see Appendix B. However, it is not clear how to use the Even-Paz protocol in this case: when the cake is connected but not convex, making parallel cuts might create disconnected pieces.

3. Fatness. A fat object is an object with a bounded length/width ratio, such as a square. Fatness makes sense in land division: if you are entitled to a 900 square meters of land, you will probably prefer to get them as a 30×30 square or a 45×20 rectangle rather than 9000×0.1 sliver. A division problem with fatness requirement cannot be reduced to one-dimensional division. In previous papers [48, 49] we developed some specialized division protocols that support fatness constraints, and they can be used instead of the Even-Paz protocol. However, we do not have an allocation-completion algorithm with fatness constraints.

4. Two pieces per agent. Theorem 1 allows an unlimited number of pieces per agent, while the other theorems allow only a single piece per agent. We do not know what happens between these two extremes. For example, what happens in a one-dimensional cake when each agent can get at most two intervals?

7.2 Improving the constants

Our redivision protocol is 1/3 or 1/4 or 1/5-proportional (depending on the geometric constraint). We see two potential ways to improve these numbers.

1. In Step #2 of our redivision protocol, we add m virtual agents, so that the total number of agents is $n+m$. However, in the Step #3, each old agent chooses either its virtual or its normal agent, while the other agent is "wasted". If we could know the m choices of the agents in advance, we could employ only n agents overall and this would subtract 1 from the constant (the constants would become 1/2 or 1/3 or 1/4). One way to analyze this scenario is to define a *strategic game* in which each old agent has two possible strategies: "normal" vs. "virtual". A pure-strategy Nash equilibrium in this game corresponds to an allocation satisfying the partial-proportionality and the gradual-retention requirements. We conjecture that a pure-strategy Nash equilibrium indeed exists. While finding a Nash equilibrium is usually a computationally-hard problem, it may be useful as an existential result.

2. In Lemma 4.2, we treat each existing piece Z_j as an "island" and insist that each new piece be entirely contained in an existing piece, i.e, we do not cross the existing division lines. This may be desirable in the context of land division, since it respects the *Uti Possidetis* principle [38]. However, it implies that the resulting division can only be partially-proportional and never fully proportional (as shown by the remark following Lemma 4.2). It may be possible to improve the proportionality guarantees by devising a different redivision procedure that does cross the existing division lines. This may require some new geometric techniques.

7.3 Price-of-fairness

It is not clear whether the upper bounds of our Theorems 7-10 are tight.

In particular, for the case of interval cake and interval pieces, there is a lower bound of $\Omega(\sqrt{n})$ on the utilitarian price of proportionality. However, we could not generalize it to the price of *partial* proportionality, and it is interesting to know which of the following two options is correct: (a) there is a lower bound of $\Omega(\sqrt{n})$ matching our Theorem 8, or (b) the actual price of partial-proportionality is $o(\sqrt{n})$. The latter option would imply that partial-proportionality is asymptotically "cheaper" than full proportionality, in social welfare terms.

APPENDIX

A Fair Division of a Rectilinear Polygon

This appendix shows what proportionality guarantees are possible when the cake is a rectilinear polygon, the pieces have to be rectangles (parallel to the sides of the cake), and there are no retention requirements. It can be seen as a baseline for Theorem 6.

Lemma A.1. *Let C be a rectilinear polygon with R reflex vertexes. It is possible to divide C among n agents such that the value of each agent is at least $1/(n+R)$ of the total cake value:*

$$\forall i \in \{1,\ldots,n\} : V_i(X_i) \geq \frac{V_i(C)}{n+R}$$

The fraction $1/(n+R)$ is the largest that can be guaranteed.

Proof. A rectilinear polygon with R reflex vertexes can be partitioned in time $O(poly(R))$ to at most $R+1$ rectangles [29, 37]. Denote these rectangles by Z_j, so that:

$$C = Z_1 \cup \cdots \cup Z_{R+1}$$

Apply the archipelago-division protocol of Lemma 4.2 with $m = R+1$. The value-guarantee per agent is at least $1/(n+m-1)$ which is at least $1/(n+R)$, as claimed.

For the upper bound, consider a staircase-shaped cake with $R+1$ stairs. as illustrated below (for $R = 4$):

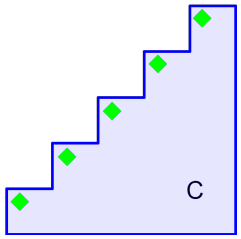

All agents have the same value-measure, which is concentrated in the diamond-shapes: the top diamond is worth n and each of the other diamonds is worth 1 (so for all agents, the total cake value is $n+R$).

Any rectangle in C can touch at most a single diamond. There are two cases:

(a) All n agents touch the top diamond. Then, their total value is n and at least one of them must receive a value of at most 1.

(b) At least one agent touches one of the R bottom diamonds. Then, the value of that agent is at most 1.

In any case, at least one agent receives at most a fraction $1/(n+R)$ of the total cake value, as claimed. □

B Cutting a Cake without Harming the Toppings

This appendix proves Lemma 5.1. For clarity, it is written in a stand-alone manner.

B.1 Introduction

Let C ("cake") be a compact subset of the Euclidean plane $\mathbb{R}^2$. Let $C_1, \ldots, C_m$ be pairwise-disjoint subsets of C ("toppings"), such that $C \supseteq C_1 \cup \cdots C_m$.

Our goal is to create a partition of C, $C = C'_1 \cup \cdots \cup C'_{m+b}$, for some $b \geq 0$, such that:

$$\forall i \in 1 \ldots m : \qquad\qquad C_i \subseteq C'_i$$

i.e, each topping is contained in a unique piece.

Subject to that requirement, it is desired to minimize b. This b is the number of empty pieces (henceforth "blanks") - the number of pieces in the final partition with no topping. What is the smallest b that can be guaranteed, in the worst case, as a function of m?

Without further restrictions on the pieces, the answer is trivial: it is always possible to get $b = 0$ by setting e.g. $C'_1 = C_1 \cup C \setminus (C_2 \cup \cdots \cup C_m)$ (arbitrarily attach all the cake outside of the toppings to one of the toppings). The question becomes more interesting when there are geometric constraints on the pieces. A geometric constraint is represented by a family S of subsets of the plane. It is given that for each $i \in \{1, \ldots, m\}$, $C_i \in S$. It is required that for each $i \in \{1, \ldots, m+b\}$, $C'_i \in S$. Several such constraints are considered here:

1. S is the family of path-connected sets, and the cake C is path-connected.
2. S is the family of convex sets, and C is convex.
3. S is the family of axes-parallel rectangles and C an axes-parallel rectangle.
4. S is the family of axes-parallel rectangles and C a rectilinear polygon.

A-priori, one could think that a sophisticated computational-geometric algorithm is needed to find a partition with a minimal b. However, I found out that the following simple algorithm provides near-optimal results:

> For each i, expand C_i arbitrarily to the largest extent possible, while keeping it contained in C and an element of S and disjoint from the other toppings.

I.e, it does not matter how exactly the toppings are expanded, to what direction and in what order, as long as they are expanded to the maximum possible amount. Formally, say that a sequence $C'_1, \ldots, C'_m$, is *maximal*, if for each $i \in 1, \ldots, m$, there is no $C^*_i \supsetneq C'_i$ such that $C^*_i \subseteq C$ and $C^*_i \in S$ and $C'_1, \ldots, C'_{i-1}, C^*_i, C'_{i+1}, \ldots, C'_m$ are pairwise-disjoint. The simple algorithm stated above can also be stated as:

(a) Find an arbitrary maximal set $C'_1, \ldots, C'_m$ such that for all i: $C'_i \subseteq C$ and $C'_i \in S$ and $C'_i \supseteq C_i$.
(b) Partition the reminder, $C \setminus (C'_1 \cup \cdots \cup C'_m)$, to a minimal number of elements of S; let that number be b and the parts $C'_{m+1}, \ldots, C'_{m+b}$.
(c) Return $C'_1, \ldots, C'_m, C'_{m+1}, \ldots, C'_{m+b}$.

Below, I analyze the performance of this algorithm for the four geometric scenarios presented above. During the analysis, for simplicity, I assume that the set $C_1, \ldots, C_m$ is already maximal. The analysis amounts to proving bounds on the number of blanks when the set of toppings is maximal.

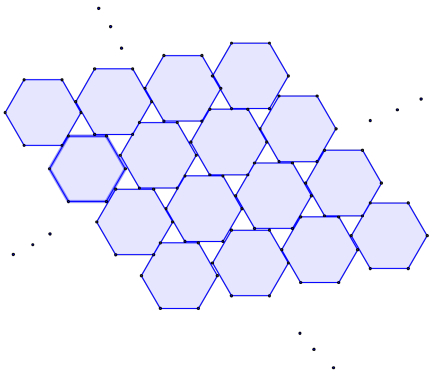

Figure 1: Lower bound on number of blanks with convex pieces.

The results are summarized in the table below. The expressions are upper and lower bounds on the worst-case number of blanks b, as a function of the number of toppings m and the number of reflex vertexes in the rectilinear polygon, R.

Cake C	Pieces S	Lower bound	Upper bound
Path-connected	Path-connected	0	0
Convex	Convex	$2m - O(\sqrt{m})$	$2m - 5$
Axes-parallel rectangle	Axes-parallel rectangles	$m - O(\sqrt{m})$	$m - 3$
Rectilinear polygon	Axes-parallel rectangles	$R + m - O(\sqrt{m})$	$R + m - 3$

B.2 Path-Connected Cake and Path-Connected Pieces

In this section, S is the family of path-connected sets and the cake C is also path-connected.

It is apparent that when the set of toppings is maximal, $b = 0$ (there are no blanks). Indeed, suppose there is a blank $H \subseteq C$. Since C is path-connected, H can be connected by a path to any of the C_i-s. In particular, there is some topping C_i such that a path from H to C_i through C does not pass through any other topping C_j. The union of C_i and the path and H is path-connected, contained in C and disjoint from the rest of the toppings, and it strictly contains C_i - in contradiction to maximality.

B.3 Convex Cake and Convex Pieces

In this section, S is the family of convex sets and the cake C is also convex. Akopyan [1] proved that the worst-case number of blanks when the set of toppings is maximal is $2m - o(m)$. I now briefly describe his proof.

B.3.1 Lower bound

Consider the tiling shown in Figure 1, and imagine that it fills the entire plane. This tiling is apparently maximal - no hexagon can be expanded while remaining convex and disjoint from the other hexagons. Each blank touches three hexagons and each hexagon touches six blanks. Hence, the number of blanks is asymptotically twice the number of hexagons.

If the tiling is finite, then the number of blanks is slightly smaller, since some blanks near the boundary of C can be discarded or joined with nearby toppings. But (assuming e.g. that C is square), the number of blanks near the boundary is only $O(\sqrt{m})$. Hence, in the worst case there are $b \geq 2m - O(\sqrt{m})$ blanks.

B.3.2 Upper bound

For the upper bound, we should distinguish between *holes* and *blanks*. A hole is a connected subset of C which is not contained in any of the toppings C_i. A-priori, a hole need not be convex. We can imagine arrangements of convex toppings, such that the space between them is not convex. Hence, the number of blanks may be much larger than the number of holes (each convex component of a hole is counted as a different blank).

However, when the set of toppings is maximal, which is the case of interest here, it can be proved that all holes are convex. Hence, in particular, each hole is a single blank, and the number of blanks equals the number of holes. The proof is in Claim 2 of [43]:

Lemma B.1. *When the set of toppings is maximal, every hole is convex.*

From here, Akopyan's proof proceeds as follows.

Each side of a hole is formed by one of the C_i-s. Lets call two sets C_i and C_j *neighbors*, if they form two adjacent sides of some hole. Note that:

- To each hole correspond at least 3 neighbor-pairs - since each hole has at least 3 sides.
- To each neighbor-pair C_i, C_j correspond at most two holes - since all such holes must have a side co-linear with the segment in which the boundaries of C_i and C_j intersect.
- Moreover, there are at least 3 neighbor-pairs near the boundary of C (since C has at least 3 sides). These neighbor-pairs correspond to at most one hole, since there is no hole in the side adjacent to the boundary.

Therefore:

$$\begin{aligned}\#\text{holes} \leq &(\#\text{pairs} - 3) \cdot 2/3 + 3 \cdot 1/3 \\ = &\#\text{pairs} \cdot 2/3 - 1\end{aligned}$$

The "neighbor" relation defines a planar graph which is denoted by G. Denote by V the set of vertexes of G - which are the toppings $C_1, \ldots, C_m$. Denote by E the set of edges of G - which are the neighbor-pairs defined above. Denote by F the set of faces of G. By Euler's formula: $|V| - |E| + |F| = 2$. Each face is at least a triangle so it is adjacent to at least three edges. Each edge is adjacent to at most two faces. Hence, $|F| \leq 2|E|/3$. Substituting into Euler's formula gives $|V| - |E|/3 \geq 2$ so $|E| \leq 3|V| - 6$. Hence, the number of neighbor-pairs is at most $3m - 6$. Hence, the number of holes is at most $2m - 5$, and so is the number of blanks: $b \leq 2m - 5$.

B.4 Rectangular Cake and Rectangular Pieces

In this section, S is the family of axes-parallel rectangles and the cake C is also an axes-parallel rectangle. I adapt the proofs of [43] and [1] to prove that the worst-case number of blanks when the set of toppings is maximal is $m - o(m)$.

B.4.1 Lower bound

Consider the tiling in Figure 2, and imagine that it fills the entire 2-dimensional plane.

This set of squares is apparently maximal. Each blank touches four squares and each square touches four blanks. Hence, the number of squares and blanks is asymptotically the same.

If the tiling is finite, then the number of blanks is slightly smaller, since some blanks near the boundary of C can be discarded or joined with nearby rectangles. But, the number of blanks near the boundary is only $\approx O(\sqrt{m})$. Hence, in the worst case there may be as many as $b \geq m - O(\sqrt{m})$ blanks.

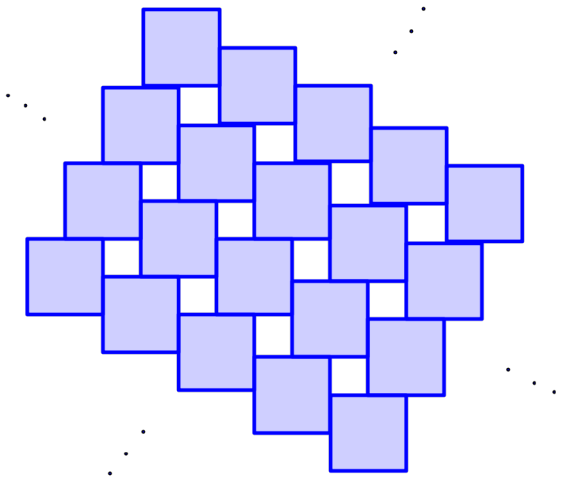

Figure 2: Lower bound on number of blanks with rectangular pieces.

B.4.2 Upper bound part A: All holes are axes-parallel rectangles

Similarly to the convex setting, I prove that every hole is an axes-parallel rectangle, and thus every hole corresponds to a single blank. Let H be a hole (a part of C not contained in any of the toppings C_i). The boundary of H is made of line-segments that are parts of the boundaries of the C_i, and possibly some line-segments that are parts of the boundary of C. The examples below refer to Figure 3.

Definition B.1. A *vertex of* H is a point on H's boundary that is a vertex of C_i or of C.

Note that this includes the vertexes of H as a polygon, but may also include "flat vertexes" (vertexes with angle 180°, such as the vertex N).

Definition B.2. An *edge of* H is a line-segment between two vertexes of H.

Definition B.3. A vertex/edge that is contained in the boundary of C is called a *boundary-vertex/edge*. Otherwise it is called an *internal-vertex/edge*.

Definition B.4. A *connection-vertex* is a boundary-vertex connecting an internal-edge and a boundary-edge (e.g. the vertexes J, O).

By our definition of the vertexes of a hole H, it is clear that each edge of H is contained either in an edge of a single topping C_i (an internal edge) or in an edge of the cake C (a boundary edge).

Definition B.5. For each edge e, let $Con(e)$ be the (unique) rectangle one of whose edges contains e. This may be either a single C_i or C.

Definition B.6. Given a vertex v and an edge e of a hole H, say that v *is exposed in* e if v is a vertex of $Con(e)$.

For example, the vertex D is exposed in the edge DQ but not in the edge DR.

Definition B.7. A vertex of H is called *reflex* if the internal angle adjacent to it (the angle through the interior of H) is larger than 180°.

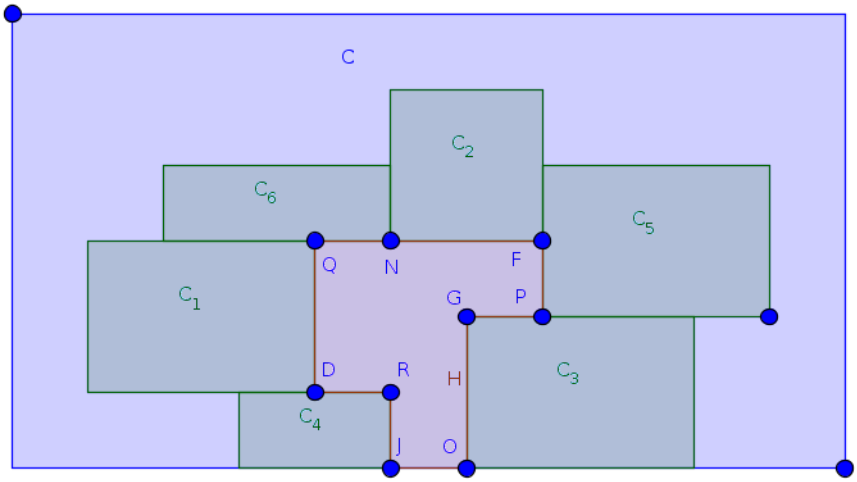

Figure 3: A rectangular cake C, rectangular toppings C_j, and a hole H. Note the configuration is not maximal.

Since all rectangles are parallel to the axes, a reflex vertex (e.g. the vertex G) must have an internal angle of 270°.

Lemma B.2. *Every internal edge of a hole cannot have both of its vertexes exposed.*

Proof. Let e be an internal edge of a hole H. Then $Con(e)$ must be one of the toppings C_i. This means that e is equal to an edge of C_i. But then, the rectangle C_i can be stretched over e and into the hole H, contradicting the maximality condition. □

For example, the internal edge DQ has both its vertexes exposed, and indeed the rectangle C_1 can be stretched rightwards into the hole and the configuration is not maximal.

Lemma B.3. *Every internal vertex is exposed in at least one of its two adjacent edges.*

Proof. Let v be an internal vertex and let e, e' be its two adjacent edges (which must both be internal edges). Suppose by contradiction that v is not exposed in e nor in e'. Then v is both internal to an edge of $Con(e) = C_i$ and internal to an edge of $Con(e') = C_j$. But this implies that C_i and C_j overlap, contradicting the pairwise-disjointness condition. □

Lemma B.4. *Every internal reflex-vertex is exposed in both its adjacent edges.*

Proof. An internal reflex-vertex of a hole is necessarily a vertex of some topping C_i which does not touch any other topping. Hence, its two adjacent edges are edges of C_i, so it is exposed in both of them (see e.g. vertexes G, R). □

Lemma B.5. *Every connection-vertex is exposed in its adjacent internal edge.*

Proof. Since the cake C is a rectangle, a connection-vertex of a hole can happen only when the vertex of a rectangle-topping touches the interior of an edge of C. □

Lemma B.6. *Every vertex of a hole must be internal.*

Proof. By contradiction, let H be a hole with a boundary-vertex. Let $e_1, \ldots, e_t$ be a maximal path of edges of H that are internal. Let v_0 be the vertex before e_1 and v_t the vertex after e_t; note that v_0 and v_t are connection-vertexes (e.g. J and O).

By Lemma B.5, v_0 is exposed in e_1 and v_t is exposed in e_t. So, in the traversal from v_0 to v_t through $e_1, \ldots, e_t$, the first vertex is exposed in the edge after it (e.g. J is exposed in JR) and

the last vertex is exposed in the edge before it (e.g. O is exposed in GO). All other vertexes in the path are internal, so by Lemma B.3 each of them is exposed in at least one edge. All in all, there are $t+1$ exposed vertexes but only t edges, so by the pigeonhole principle, at least one edge in the path must have two of its vertexes exposed - a contradiction to Lemma B.2. □

Lemma B.7. *No vertex of a hole H is reflex.*

Proof. By contradiction, let v be a reflex vertex of H (e.g. the vertex G). By Lemma B.6 it is internal so by Lemma B.4 it is exposed in both its adjacent edges.

Starting from v, traverse the edges and vertexes of H until you reach v again (e.g. from G through GP to P through PF to F ... to O through OG to G). In this traversal, v is exposed both in the edge after it (e.g. GP) and in the edge before it (e.g. GO). By Lemma B.6, all vertexes in the traversal are internal, so by Lemma B.3, each of them is exposed in at least one edge. All in all, there are $t+1$ "exposed" relations but only t edges, so by the pigeonhole principle, at least one edge in the path must have two of its vertexes exposed - a contradiction to Lemma B.2. □

Corollary B.1. *Every hole H must be a rectangle and must be strictly in the interior of C (not adjacent to its boundary).*

B.4.3 Upper bound part B: Counting the holes

By Corollary B.1, each hole must be a rectangle internal to the cake C. Therefore, the sides of each hole are formed by the rectangular toppings $C_1, \ldots, C_m$. Let's call two such toppings *neighbors* if they form adjacent sides of the same hole.

Each hole has 4 sides, so to each hole correspond (at least) 4 neighbor-pairs.

To each neighbor-pair C_i, C_j correspond at most two holes - since all such holes must have a side co-linear with the segment in which the boundaries of C_i and C_j intersect.

Moreover, there are at least 4 neighbor-pairs near the boundary of C (since C has 4 sides). These neighbor-pairs correspond to at most one hole, since there is no hole in the side adjacent to the boundary.

Therefore:

$$\begin{aligned}\#\text{holes} \leq &(\#\text{pairs} - 4) \cdot 2/4 + 4 \cdot 1/4 \\ = &\#\text{pairs}/2 - 1\end{aligned}$$

The "neighbor" relation defines a planar graph which is denoted by G. Denote by V the set of vertexes of G - which are the rectangles $C_1, \ldots, C_m$. Denote by E the set of edges of G - which are the neighbor-pairs defined above. Denote by F the set of faces of G. By Euler's formula: $|V| - |E| + |F| = 2$. Since all angles are multiples of 90°, each face is adjacent to at least four edges. Each edge is adjacent to at most two faces. Hence, $|F| \leq |E|/2$. Substituting into Euler's formula gives $|V| - |E|/2 \geq 2$ so $|E| \leq 2|V| - 4$. Hence, the number of neighbor-pairs is at most $2m - 4$, so the number of holes is at most $m - 3$ and so is the number of blanks: $b \leq m - 3$.

B.5 Rectilinear-polygonal cake and rectangular pieces

In this section, S is still the family of axes-parallel rectangles, but now the cake C can be an arbitrary axes-parallel rectilinear polygon. The "complexity" of a rectilinear polygon is characterized by the number of its *reflex vertexes* - vertexes with internal angle 270°. It is known that a rectilinear polygon with R reflex vertexes can always be partitioned to at most $R+1$ rectangles [29, 37], and this bound is tight when the vertexes of C are in general position. Since our goal is to bound b - the number of blank rectangles - it is expected that the bound should depend on R, in addition to m (the number of toppings).

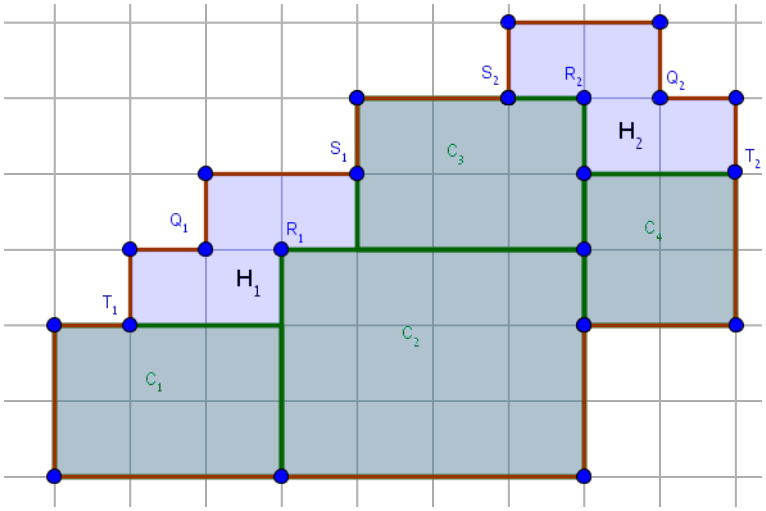

Figure 4: A rectilinear cake with four rectangular topping $C_1, \ldots, C_4$, and two rectilinear holes H_1, H_2. Note the configuration is not maximal since C_4 can be extended upwards into H_2.

B.5.1 Lower bound

Consider the tiling of Figure 2 in Subsection B.4.1. Assume that the tiling is contained in a very large square, and the square is connected in a narrow pipe to a rectilinear polygon with R reflex vertexes. As explained there, there are $m - O(\sqrt{m})$ blanks in the square. The connected rectilinear polygon is a single *hole*, but it may contain as many as $R+1$ *blanks* (its rectangular components). All in all, there are $m + R - O(\sqrt{m})$ blanks.

B.5.2 Upper bound part A: characterizing the holes

The first three lemmas of Subsection B.4.2 (Lemmas B.2 and B.3 and B.4) are still true since they do not depend on the shape of C. But the last three lemmas are not true.

The difference starts at Lemma B.5: since C is not a rectangle, there may be connection-vertexes which are vertexes of C and not of any topping C_i. See e.g. the vertexes S_1, T_1, S_2 in Figure 4. That lemma has to be generalized as follows:

Lemma B.8. *Every connection-vertex of a hole is either exposed in its adjacent internal edge, or it is a reflex vertex of the cake C.*

Proof. Let v be a connection-vertex. If it is a vertex of a topping, then it is exposed in the edge of that topping (which is an internal edge), so we are done.

Otherwise, it must be a vertex of the cake. If it is a convex vertex of the cake, then it necessarily connects two edges of the cake, so it cannot be a connection-vertex. Hence, it must be a reflex vertex of the cake. □

Note that, if a connection-vertex of a hole H is a reflex-vertex of C (the latter case in the above lemma), then that vertex must be a convex vertex of H. Hence, such a vertex is called a *convexified vertex*:

Definition B.8. A *convexified vertex* is a reflex-vertex of C which is a convex-vertex of some hole.

So by Lemma B.8, every connection-vertex of a hole is either exposed in its adjacent internal edge, or it is a convexified vertex.

Lemma B.6 is also no longer true; a hole may be adjacent to the boundary of C. That lemma has to be generalized as follows:

Lemma B.9. *If a hole is adjacent to the boundary of C, then it must have at least one convexified vertex.*

Proof. Each hole must have at least two connection-vertexes. If, by contradiction, none of these vertexes is convexified, then by Lemma B.8, both must be exposed in the internal edges adjacent to them. But then, similarly to the proof of Lemma B.6, the path between these vertexes must have an edge both of whose vertexes are exposed - a contradiction. □

Lemma B.7 is also no longer true. A hole may have reflex vertexes. E.g, in H_1, both Q_1 and R_1 are reflex vertexes; Q_1 is a reflex boundary-vertex (it is a reflex vertex of C), and R_1 is a reflex inner-vertex. That lemma has to be generalized as follows:

Lemma B.10. *Let v be a reflex internal-vertex of a hole H. Then, any traversal of the boundary of H starting at v, either clockwise or counter-clockwise, must eventually meet a convexified vertex of C, and it must not meet any other reflex vertex of H on its way.*

Proof. Lemma B.4 is still true, so regardless of the direction of our traversal, v is exposed in the edge following it in the path. If it meets another reflex vertex (including itself), then, since this another vertex is exposed in the edge before it in the path, by the pigeonhole principle there must be an edge both of whose vertexes are exposed - a contradiction (this reasoning is similar to the proof of Lemma B.7). Similarly, if it meets a connection-vertex which is not convexified (like T_2), then this vertex is exposed in the edge before it in the path, which again leads to a contradiction. The only remaining option is that the traversal ends in a convexified vertex, since this is the only kind of vertex which may be not exposed. □

For example, a traversal from R_1 clockwise meets T_1 and a traversal counter-clockwise meets S_1; both are convexified vertexes of C, and both are not exposed in the edges between them and R_1.

The above lemma has the following corollary, which replaces Corollary B.1:

Corollary B.2. *Every internal hole (a hole not adjacent to the boundary of C) must be a rectangle.*

B.5.3 Upper bound part B: Counting the holes

The upper bound in Subsection B.4.3, namely $m-2$, is still a valid upper bound on the number of internal holes. By Corollary B.2, this is a valid bound on the number of internal blanks.

To bound the total number of blanks, we have to handle the blanks resulting from boundary holes. The main step is to bound the number of reflex vertexes in the boundary holes.

Lemma B.11. *Let B_h be the set of boundary holes, R_h the set of reflex vertexes in all boundary holes, and R_C the set of reflex vertexes in C. Then:*

$$|B_h| + |R_h| \leq |R_C|$$

Proof. The lemma is proved by the following 1-1 mapping from $B_h \cup R_h$ into R_C:

- Each element of B_h (a boundary hole) is mapped to a convexified vertex of C, which exists by Lemma B.8.
- Each element $v \in R_h$ which is an internal reflex-vertex, by Lemma B.10 corresponds to two convexified-vertexes; at most one of them is already mapped to the hole containing v, so v can be mapped to the other convexified-vertex.
- Each element $v \in R_h$ which is a boundary reflex-vertex is mapped to itself. Note that v is a reflex vertex of a hole, so it is not convexified, so it was not mapped in the previous two bullets.

□

Corollary B.3. *The total number of reflex-vertexes in all boundary holes satisfies:*

$$|R_h| \leq R - |B_h|$$

where $R = |R_C| =$ *the number of reflex vertexes of* C.

As mentioned above, a rectilinear polygon with R reflex vertexes can always be partitioned to at most $R + 1$ rectangles [29, 37]. Hence, by Corollary B.3, the set B_h can be partitioned to rectangles such that the total number of rectangles is at most $(R - |B_h|) + |B_h| = R$.

All in all, the number of (rectangular) blanks is at most: $R + m - 3$.

References

[1] A. Akopyan. Partitioning a convex object without harming existing convex subsets. Math Overflow, Feb. 2016. URL:http://mathoverflow.net/a/232445/34461 (version: 2016-03-14).

[2] O. Arzi. Cake Cutting: Achieving Efficiency While Maintaining Fairness. Master's thesis, Bar-Ilan University, Oct. 2012. Under the supervision of Prof. Yonatan Aumann.

[3] Y. Aumann and Y. Dombb. The Efficiency of Fair Division with Connected Pieces. *Web, Internet and Network Economics*, 6484:26–37, 2010.

[4] Y. Aumann, Y. Dombb, and A. Hassidim. Computing Socially-Efficient Cake Divisions. In *The International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 343–350, May 2013.

[5] H. Aziz and C. Ye. Cake Cutting Algorithms for Piecewise Constant and Piecewise Uniform Valuations. In T.-Y. Liu, Q. Qi, and Y. Ye, editors, *Web and Internet Economics*, volume 8877 of *Lecture Notes in Computer Science*, pages 1–14. Springer International Publishing, 2014.

[6] J. B. Barbanel, S. J. Brams, and W. Stromquist. Cutting a pie is not a piece of cake. *American Mathematical Monthly*, 116(6):496–514, 2009.

[7] A. Beck. Constructing a Fair Border. *The American Mathematical Monthly*, 94(2):157–162, Feb. 1987.

[8] X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal Proportional Cake Cutting with Connected Pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1263–1269, 2012.

[9] M. Berliant and K. Dunz. A foundation of location theory: existence of equilibrium, the welfare theorems, and core. *Journal of Mathematical Economics*, 40(5):593–618, Aug. 2004.

[10] M. Berliant, W. Thomson, and K. Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, Jan. 1992.

[11] H. Bernstein. Land Reform: Taking a Long(er) View. *Journal of Agrarian Change*, 2(4):433–463, Oct. 2002.

[12] D. Bertsimas, V. F. Farias, and N. Trichakis. The Price of Fairness. *Operations Research*, 59(1):17–31, Feb. 2011.

[13] D. Bertsimas, V. F. Farias, and N. Trichakis. On the Efficiency-Fairness Trade-off. *Manage. Sci.*, 58(12):2234–2250, Dec. 2012.

[14] S. J. Brams. *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press, first edition edition, Dec. 2007.

[15] S. J. Brams, M. Feldman, J. K. Lai, J. Morgenstern, and A. D. Procaccia. On Maxsum Fair Cake Divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1285–1291, 2012.

[16] S. J. Brams, M. Jones, and C. Klamler. Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4):353–367, Mar. 2008.

[17] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, Feb. 1996.

[18] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The Efficiency of Fair Division. *Theory of Computing Systems*, 50(4):589–610, Sept. 2012.

[19] I. Caragiannis, J. K. Lai, and A. D. Procaccia. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 127–132. AAAI Press, 2011.

[20] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30(1):3–31, 2006.

[21] Y. J. Cohler, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Optimal Envy-Free Cake Cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI-11)*, pages 626–631, 2011.

[22] M. Dall'Aglio and F. Maccheroni. Disputed lands. *Games and Economic Behavior*, 66(1):57–77, May 2009.

[23] C. Delimitrou and C. Kozyrakis. Quasar: Resource-efficient and QoS-aware Cluster Management. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '14, pages 127–144, New York, NY, USA, 2014. ACM.

[24] R. Devulapalli. *Geometric partitioning algorithms for fair division of geographic resources.* PhD thesis, University of Minnesota, July 2014.

[25] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of Fairness in Kidney Exchange. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 1013–1020, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[26] L. E. Dubins and E. H. Spanier. How to Cut A Cake Fairly. *The American Mathematical Monthly*, 68(1):1+, Jan. 1961.

[27] J. Edmonds and K. Pruhs. Balanced Allocations of Cake. In *FOCS*, volume 47, pages 623–634. IEEE Computer Society, Oct. 2006.

[28] J. Edmonds, K. Pruhs, and J. Solanki. Confidently Cutting a Cake into Approximately Fair Pieces. *Algorithmic Aspects in Information and Management*, pages 155–164, 2008.

[29] D. Eppstein. Graph-Theoretic Solutions to Computational Geometry Problems. In C. Paul and M. Habib, editors, *Graph-Theoretic Concepts in Computer Science*, volume 5911 of *Lecture Notes in Computer Science*, pages 1–16. Springer Berlin Heidelberg, 2010.

[30] S. Even and A. Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296, Mar. 1984.

[31] E. Friedman, C. A. Psomas, and S. Vardi. Dynamic Fair Division with Minimal Disruptions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 697–713, New York, NY, USA, 2015. ACM.

[32] T. P. Hill. Determining a Fair Border. *The American Mathematical Monthly*, 90(7):438–442, Aug. 1983.

[33] M. Hohenwarter, M. Borcherds, G. Ancsin, B. Bencze, M. Blossier, A. Delobelle, C. Denizet, J. Éliás, A. Fekete, L. Gál, Z. Konečný, Z. Kovács, S. Lizelfelner, B. Parisse, and G. Sturr. GeoGebra 5.0, Dec. 2013.
urlhttp://www.geogebra.org.

[34] T. Ichiishi and A. Idzik. Equitable allocation of divisible goods. *Journal of Mathematical Economics*, 32(4):389–400, Dec. 1999.

[35] K. Iyer and M. N. Huhns. A Procedure for the Allocation of Two-Dimensional Resources in a Multiagent System. *International Journal of Cooperative Information Systems*, 18:1–34, 2009.

[36] I. Kash, A. D. Procaccia, and N. Shah. No agent left behind: dynamic fair division of multiple resources. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, AAMAS '13, pages 351–358. IFAAMAS, 2013.

[37] J. M. Keil. Polygon Decomposition. In *Handbook of Computational Geometry*, pages 491–518. University of Saskatchewan Saskatoon, Sask., Canada, 2000.

[38] S. N. Lalonde. *Determining boundaries in a conflicted world: the role of Uti Possidetis*. McGill-Queen's Press-MQUP, 2002.

[39] M. Lipton. *Land Reform in Developing Countries: Property Rights and Property Wrongs*. Routledge, Aug. 2009.

[40] K. McAvaney, J. M. Robertson, and W. A. Webb. Ramsey partitions of integers and fair divisions. *Combinatorica*, 12(2):193–201, 1992.

[41] H. Moulin. Fair division under joint ownership: Recent results and open problems. *Social Choice and Welfare*, 7(2):149–170, 1990.

[42] H. Moulin. *Fair Division and Collective Welfare*. The MIT Press, Aug. 2004.

[43] R. Pinchasi. On the perimeter of k pairwise disjoint convex bodies contained in a convex set in the plane. *Combinatorica*, pages 1–27, July 2015.

[44] J. P. Powelson. *The Story of Land: A World History of Land Tenure and Agrarian Reform*. Lincoln Institute of Land Policy, May 1988.

[45] A. D. Procaccia. Cake Cutting Algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2015.

[46] J. M. Robertson and W. A. Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. A K Peters/CRC Press, first edition, July 1998.

[47] P. Rosset, R. Patel, and M. Courville, editors. *Promised Land: Competing Visions of Agrarian Reform*. Food First Books, first trade paper edition edition, Nov. 2006.

[48] E. Segal-Halevi, A. Hassidim, and Y. Aumann. Envy-Free Cake-Cutting in Two Dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 1021–1028, Jan. 2015.

[49] E. Segal-Halevi, S. Nitzan, A. Hassidim, and Y. Aumann. Fair and Square: Cake-Cutting in Two Dimensions, Oct. 2015. arXiv preprint 1510.03170.

[50] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, Jan. 1948.

[51] W. Thomson. The Fair Division of a Fixed Supply Among a Growing Population. *Mathematics of Operations Research*, 8(3):319–326, Aug. 1983.

[52] W. Thomson. Children Crying at Birthday Parties. Why? *Economic Theory*, 31(3):501–521, June 2007.

[53] W. Thomson. *Fair Allocation Rules*, volume 2, pages 393–506. Elsevier, 2011.

[54] T. Walsh. Online Cake Cutting. *Algorithmic Decision Theory*, 6992:292–305, 2011.

[55] Y. Zhang, M. A. Laurenzano, J. Mars, and L. Tang. SMiTe: Precise QoS Prediction on Real-System SMT Processors to Improve Utilization in Warehouse Scale Computers. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-47, pages 406–418, Washington, DC, USA, 2014. IEEE Computer Society.

[56] R. Zivan. Can trust increase the efficiency of cake cutting algorithms. In *The International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1145–1146, 2011.